A Delay Efficient Multiclass Packet Scheduler for Heterogeneous M2M Uplink

Akshay Kumar, Ahmed Abdelhadi and Charles Clancy
Hume Center, Virginia Tech
Email:{akshay2, aabdelhadi, tcc}@vt.edu

arXiv:1601.03061v2 [cs.NI] 23 Apr 2016

***Abstract*—The sensory traffic in Machine-to-Machine (M2M) communications has fairly heterogeneous service delay requirements. Therefore, we study the delay-performance of a heterogeneous M2M uplink from the sensors to a M2M application server (AS) via M2M aggregators (MA). We classify the heterogeneous M2M traffic aggregated at AS into multiple Periodic Update (PU) and Event Driven (ED) classes. The PU arrivals are periodic and need to be processed by a prespecified *firm* service deadline whereas the ED arrivals are random with *firm* or *soft* real-time or non real-time service requirements. We use step and sigmoidal functions to represent the service utility for PU and ED packets respectively. We propose a delay efficient multiclass packet scheduling heuristic that aims to maximize a proportionally fair system utility metric. Specifically, the proposed scheduler prioritizes service to ED data while ensuring that the PU packets meet their service deadline. It also minimizes successive PU failures for critical applications by penalizing their occurrences. Furthermore, the failed PU packets are immediately cleared from the system so as to reduce network congestion. Using extensive simulations, we show that the proposed scheduler outperforms popular packet schedulers and the performance gap increases with heterogeneity in latency requirements and with greater penalty for PU failures in critical applications.**

***Index Terms*—M2M, Latency, Quality-of-Service, Multiclass Scheduler**

I. Introduction

Machine-to-Machine (M2M) communications is becoming increasingly ubiquitous due to its vast number of residential and industrial applications such as in smart homes, vehicle tracking, industrial automation etc. In most of the M2M applications, the uplink traffic generated from the sensors is fairly heterogeneous and can be classified as either non real-time (no deadline for task completion) or *soft* real-time (decreased utility if deadline not met) or *firm* (zero utility if deadline not met). For instance, the messages from a instrumented protective system have *firm* real-time service requirements whereas preventive maintenance applications are typically served in non real-time. With increase in network size, the available computational and communication resources gets shared among a large number of sensors which makes it very hard to provide real-time service [1]. This necessitates the need of designing delay-efficient packet scheduler for heterogeneous M2M uplink and that is applicable for any M2M application, communication standard and hardware-software architecture.

This research is based upon work supported by the National Science Foundation under Grant No. 1134843.

Most of the existing M2M packet schedulers are designed for specific wireless technology such as LTE (see [2] and references therein) and fail to fully characterize the heterogeneity in M2M uplink traffic as done in our work. Another line of work considers the design of packet schedulers for real-time control applications (see [3] and references therein) that guarantee the schedulability of critical tasks while providing best-effort service for others. These works assume that the relative deadline for periodic tasks is equal to their periods, all aperiodic tasks are aggregated into a single class and the minimum inter-arrival time of an aperiodic task equals its deadline. In our work, we do not make any of these restrictive assumptions on the nature of M2M traffic. Also, unlike these works, the objective of our scheduler is to ensure utility-proportional fairness among all traffic classes.

The work that is most closely related to our work is by Kumar et. al. in [4] wherein a delay-efficient heuristic scheduler is proposed for M2M uplink. However, this work does not account for the diverse delay requirements of traffic from different sensors and it classifies the aggregated sensor traffic into two broad PU and ED classes. Also, it does not attempts to minimize successive PU failures, which may result in undesired consequences for the M2M system.

In this paper, we develop a delay-efficient mutliclass packet scheduler for heterogeneous M2M uplink traffic. To begin with, we use source M2M traffic model in [5] to classify the data from each sensor as Periodic Update (PU) or Event Driven (ED). The PU arrivals are periodic with *firm* service deadline while the ED arrivals are random with *firm* or *soft* real-time or non real-time service requirements. We map the contrasting delay requirements of PU and ED packets onto step and sigmoidal utility functions respectively. The goal of proposed scheduler is to maximally satisfy the delay requirement of all classes which translates to maximizing a proportionally-fair system utility metric. It does so by prioritizing service to ED packets while ensuring that the PU packets meet their deadline. This results in a higher utility for the ED packets without compromising the utility for PU packets.

However, we note that, due to random arrival and service times, some PU packets fail to meet their deadline. It becomes particularly difficult to do so as congestion increases with increase in network size. When this occurs, we remove failed PU packets to reduce network congestion and improve overall system utility. Furthermore, for critical M2M applications,

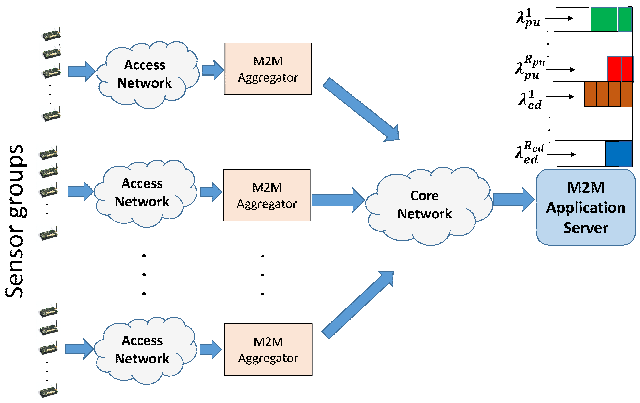

Fig. 1. System model illustrating queuing process at M2M AS.

successive PU failures are highly detrimental to system performance. Hence, we aim to minimize successive PU failures by imposing a penalty to the PU utility. Using extensive simulations, we show that the proposed scheduler has a superior delay-performance compared to popular scheduling policies such as First-Come-First-Serve (FCFS), Earliest Due Date (EDD), fixed priority scheduler etc. The performance gap increases with heterogeneity in latency requirements and with higher penalty for successive PU failures for critical applications. Another salient feature of the proposed scheduler is that it is general enough to be used across different M2M applications, M2M communication standards and hardware-software M2M architectures.

The rest of the paper is organized as follows. Section II introduces the system model and defines the utility functions for PU and ED data. Then in Section III, we formulate the utility maximization problem and in Section IV describe the proposed packet scheduling heuristic. Section V presents simulation results. Finally Section VI draws some conclusions.

II. System Model

Fig 1 shows the system model for a heterogeneous M2M uplink from N sensors to a M2M application server (AS) via multiple M2M aggregators. The aggregated M2M traffic at AS is classified into R_{pu} PU and R_{ed} ED classes. Here, we assume that there exists multiple sensors with similar arrival period and service deadlines for PU packets and therefore the PU traffic at AS can be clustered into R_{pu} classes. Similarly we cluster the aggregated ED traffic into R_{ed} classes based on the closeness in the characteristics of ED traffic. The period and relative deadline for i^{th} PU class are T^i_{pu} and l^i_d respectively. The ED arrivals is modeled as a Poisson process with rate λ^i_{ed} for i^{th} class.

The service times for i^{th} PU and ED class are assumed to be exponential with rates μ^i_{pu} and μ^i_{ed} respectively. In this work, we consider only queuing and service delay at AS and ignore other components. Due to relatively small payload for M2M traffic, the packet transmission time are small relative to queuing delay and thus can be ignored. We assume sufficient number of wireless resources for both access and core network to ignore the congestion delay. Now the scheduling policy at AS should be chosen as to maximally satisfy the delay requirements of each PU and ED class.

III. Problem Formulation

A. Utility Functions

1) PU utility: As stated previously, the PU packets need to be processed within a prespecified time interval at the end of which there is no utility in serving the packet. For a PU packet of i^{th} class with latency l^i_{pu}, we define the utility function (see Fig. 2) as,

$$U^i_{\text{pu}}(l^i_{\text{pu}}) = \begin{cases} 1 & \text{if } l^i_{pu} < l^i_d \\ 0 & \text{if } l^i_{pu} \geq l^i_d, \end{cases} \tag{1}$$

where l^i_d is the latency deadline at which utility drops to 0.

We next define a penalty function to limit the number of successive PU failures. Let r_i denote the run-length[1] of a sequence of PU failures for i^{th} class, then the penalty for the run (besides the utility loss for each PU failure) is,

$$P^i_{\text{pu}}(r_i) = r_i - {r_i}^{\gamma_i}, \; r_i \geq 1 \tag{2}$$

where $\gamma_i \geq 1$ indicates the severity of the penalty for i^{th} class and is chosen based on the criticality of the application. For instance, consider the smart-grid uplink traffic which besides other messages, consists of periodic real-time pricing (RTP) updates and periodic meter readings [6]. The RTP updates are highly critical messages and failure to process them within a particular time can potentially result in significant economic loss [7]. On the contrary, failure to process periodic meter readings in time is not a big concern. It can easily be estimated using the meter reading of last successfully received (at the M2M application server) meter reading and any differences in the estimate and actual reading can be adjusted in future billing cycle. Therefore, the penalty (γ_i) for RTP failures should be higher than that for meter readings. The special case of $\gamma_i = 1$ corresponds to when there is no penalty for successive PU failures.

2) ED utility: Based on the M2M application, the delay requirements of ED traffic may vary from *firm* to non real-time service requirements. Sigmodial functions are fairly versatile to model diverse Quality-of-Service requirements by appropriately varying their parameters [8], [9]. Therefore, we use sigmoidal function to model the utility of i^{th} ED class (see Fig. 3) as,

$$U^i_{\text{ed}}(l^i_{\text{ed}}) = 1 - c_i \left(\frac{1}{1 + \text{e}^{-a_i(l^i_{\text{ed}} - b_i)}} - d_i \right), \tag{3}$$

where, $c_i = \frac{1+\text{e}^{a_i b_i}}{\text{e}^{a_i b_i}}$ and $d_i = \frac{1}{1+\text{e}^{a_i b_i}}$. The parameter a_i is the utility roll-off factor that signifies the criticality of the application while the parameter b_i roughly indicates a soft latency deadline.

[1]A run is a maximal length sequence of identical outcomes within a larger sequence. For example, let $[SSFFSFSFFFS]$ denote a sequence of PU success (S) and failure (F). Here, we have three runs of PU failures with length r_i equal to 1, 2 and 3 respectively.

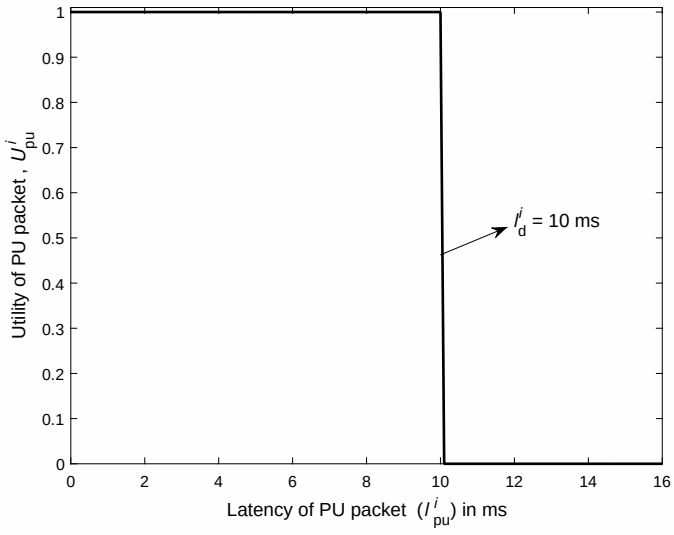

Fig. 2. Utility function for PU packet of class i.

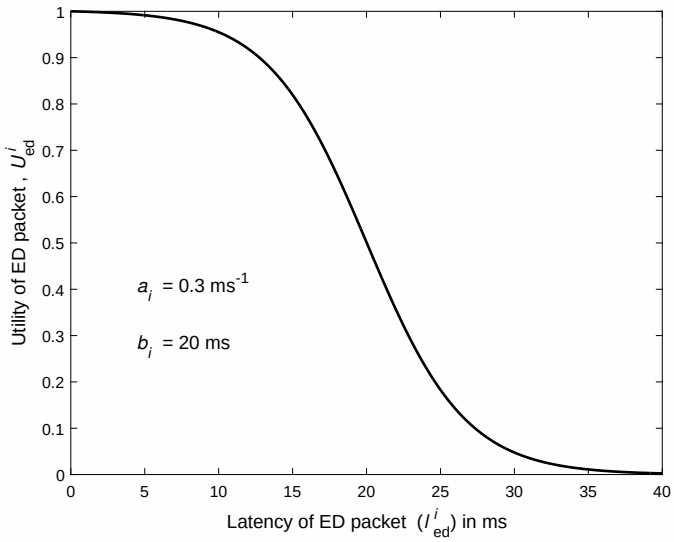

Fig. 3. Utility function for ED packet of class i.

B. System utility function

For a given scheduling policy $\mathcal{P}$, let $\hat{U}^i_{\text{pu}}(\mathcal{P})$ and $\hat{U}^i_{\text{ed}}(\mathcal{P})$ denote the average utility of i^{th} PU and ED class respectively, as given by,

$$\hat{U}^i_{\text{pu}}(\mathcal{P}) = \lim_{T_s \to \infty} \frac{\sum_{j=1}^{M^i_{\text{pu}}(T_s)} U^i_{\text{pu}}(l^{ij}_{\text{pu}}) + \sum_{j=1}^{R_i(T_s)} P^i_{\text{pu}}(r_{ij})}{M^i_{\text{pu}}(T_s)}, \quad (4)$$

$$\hat{U}^i_{\text{ed}}(\mathcal{P}) = \lim_{T_s \to \infty} \frac{\sum_{j=1}^{M^i_{\text{ed}}(T_s)} U^i_{\text{ed}}(l^{ij}_{\text{ed}})}{M^i_{\text{ed}}(T_s)}, \quad (5)$$

where $M^i_{\text{pu}}(T_s)$ and $M^i_{\text{ed}}(T_s)$ are the number of packets from i^{th} PU and ED class served in time T_s. The variable r_{ij} denotes the length of j^{th} sequence or run of PU failures for class i and $R_i(T_s)$ denotes the number of PU failure sequences in time T_s. Lastly, l^{ij}_{pu} and l^{ij}_{ed} denote the latency of the j^{th} packet i^{th} PU and ED class respectively.

We next aggregate the average PU and ED utilities into the following proportionally-fair system utility metric.

$$V(\mathcal{P}) = \prod_{i=1}^{R_{\text{pu}}} (\hat{U}^i_{\text{pu}})^{\beta^i_{\text{pu}}}(\mathcal{P}) * \prod_{i=1}^{R_{\text{ed}}} (\hat{U}^i_{\text{ed}})^{\beta^i_{\text{ed}}}(\mathcal{P}), \quad (6)$$

where the parameters β^i_{pu} and β^i_{ed} denote the relative priority assigned to i^{th} PU and ED class respectively.

C. Delay-optimal scheduling policy

Our objective is to determine the optimal scheduling policy $\mathcal{P}$ at M2M AS that maximizes the proposed system utility metric, $V(\mathcal{P})$. If the arrival time and service times of packets of each PU and ED class are known *apriori*, then the optimal scheme is to process packets of i^{th} PU class such that their service is completed exactly at their deadline. However, due to random ED arrivals and random service times, it is not possible to determine an online optimal scheduler [10]. Therefore, we propose an online heuristic scheduler with a goal to maximize system utility, as described in the next section.

IV. Proposed Scheduler

The proposed scheduler dynamically schedules packets of different classes using delay-efficient heuristics that aim to maximize the proposed system utility metric. Therefore, the backlogged packets at the AS may preempt the packet in service depending on the following scheduling heuristics.

A. Service order between PU and ED classes

From the PU utility function in Fig. 2, it is clear that there is no added utility of processing it before the deadline l^i_d. On the contrary, the ED utility decreases monotonically with increase in latency. Therefore, the proposed scheduler should prioritize service to backlogged ED packets until the delay for PU packet of class i exceeds certain delay threshold $l^i_t (l^i_t < l^i_d)$. Note that, due to randomness in arrival and service times of packets, it is not possible to guarantee that all PU packets will meet their deadline. However, the value of l^i_t is chosen such that the .

B. Service order among PU classes

For resolving service contention between packets of different PU classes, we give higher (preemptive) priority to the PU class with higher penalty factor γ_i and lower service rate μ^i_{pu}. This is because higher γ_i for PU class indicates that it is less robust to PU failures and therefore, should be served first. We prioritize service to PU class with lower service rate (i.e., higher average service time) because it is more likely to be unsuccessful if its service is delayed. Again due to the *firm* utility function for PU classes, if there are backlogged ED packets, then the priority order between two PU classes comes into effect only when packets of each PU class exceeds its corresponding latency threshold.

C. Service order among ED classes

We now discuss how the proposed scheduler determines the order of service between packets of different ED classes. We assign higher priority to ED classes that are more delay sensitive (i.e., have higher a_i and/or lower b_i). Unlike the earlier heuristics, the (preemptive) priority order among ED classes comes into effect immediately at the arrival of ED packet. However, if the latency of i^{th} ED class exceeds a certain threshold δ_i while there are backlogged packet of lower priority ED classes, then its gets preempted by the ED class with the highest priority among lower priority (backlogged) ED classes. By definition, the threshold $\delta_i = \infty$ for ED class with least priority.

The preemption policy for ED classes exceeding their threshold δ_i ensures that ED packets with unusually large service time do not hog the server which will greatly increase the latency for other ED classes. Thus, the proposed heuristic aims to reduce the latency of delay-sensitive ED class without incurring very large latency for delay-tolerant classes.

D. Service order among packets of same class

We serve the packets of a given PU or ED class using FCFS policy. Although the Shortest Processing Time (SPT) algorithm is optimal for minimizing the average latency [11], we cannot implement it here because we assume random service time for packets and hence this information is not known *apriori* to the scheduler. Since the average service time is same for all packets of a particular class, FCFS is the best policy in this case.

E. Complexity Analysis

For the proposed scheduler to maximize the system utility, we need to determine the jointly optimal value of PU and ED thresholds i.e., l_t^i, δ_i. However, the computational complexity for this optimization may be huge if we perform a brute force search over all permissible values of thresholds l_t^i and δ_i, particularly for large number of PU and ED classes. Therefore, we next present the heuristics to reduce the search space for optimization, without sacrificing much on the delay-performance.

Search space for the PU threshold: The unconstrained search space for the PU threshold for class i is $l_t^i \in [0, l_d^i]$. We have assumed the service time for PU packets of class i to be exponentially distributed with rate μ_{pu}^i. So, the 99 percentile for service time is roughly $t_{99}^i = 4.6/\mu_{\text{pu}}^i$. Therefore, we reduce the search space for class i to $l_t^i \in [l_d^i - t_{99}^i, l_d^i]$ without significantly affecting its delay-performance. The computational savings are significant for PU class with large deadline (l_d^i) and high service rate (μ_{pu}^i).

Search space for the ED threshold: The unconstrained search space for the i^{th} ED class threshold is $\delta_t^i \in [0, \infty]$. Let $l_{\text{ed},x}^i$ denote the latency at which the utility of packet of i^{th} ED class equals x and can be easily determined using (3). We set the upper bound on δ_t^i to $l_{\text{ed},0.01}^i$. Since a class i packet on average takes $t_{\text{av}}^i = 1/\mu_{\text{pu}}^i$ time at server, a reasonable lower bound for δ_t^i is $\min(0, l_{\text{ed},0.99}^i - t_{\text{av}}^i)$. Hence, the reduced search space for class i is $\delta_t^i \in [\min(0, l_{\text{ed},0.99}^i - t_{\text{av}}^i), l_{\text{ed},0.01}^i]$.

Lastly, the proposed scheduler drops the backlogged PU packets that have exceeded their latency deadline as there utility drops to 0. However, this helps clear congestion and thus reduce latency for other backlogged packets. The proposed scheduler is described in detail in Algorithm 1.

Algorithm 1 Proposed Mutliclass Packet Scheduler

Begin Algorithm:

while there are backlogged packets **do**
  Remove the failed PU packets.
  Find classes with packet delay exceeding threshold.
  if some PU classes exceed their threshold **then**
    Serve the PU class with highest priority among PU class that exceed their threshold.
  else if no PU or ED class exceed its threshold **then**
    Serve the ED class with overall highest priority.
  else if no PU but some ED classes exceed their threshold **then**
    Serve the ED class with highest priority among ED classes that do not exceed their threshold.
  else
    Continue processing the current packet in server.
  end if
end while

V. Simulation Results

In this section, we use Monte-Carlo simulations to evaluate the system utility performance of our scheduler against various standard scheduling policies such as FCFS, EDD, Preemptive priority scheduling. The simulation time T_s is set to a large value (40 s) to ensure a steady-state queuing behavior. Although, the proposed scheduler is designed for any value of R_{pu} and R_{ed}, we set $R_{\text{pu}} = R_{\text{ed}} = 2$ to gain practical insights about its performance. The total number of sensors is $N = 500$. Of these, 300 sensors transmit packets of PU class 1 and the remaining 200 transmit packets of PU class 2. For ED traffic, we assume 150 sensors for class 1 and the remaining 350 for class 2. The PU period for each sensor of type class 1 and 2 are $T_{\text{pu}}^1 = 1.8$ s and $T_{\text{pu}}^2 = 0.5$ s respectively. The overall arrival rate for ED class 2 is $\lambda_{\text{ed}}^2 = 0.1$ packet/ms whereas λ_{ed}^1 is varied from 0.1–0.25 packet/ms. The service rates for each class is set to 1 packet/ms.

All classes are assigned equal priority, unless mentioned otherwise. The latency deadline for PU class 1 and 2 are $l_d^1 = 4$ ms and $l_d^2 = 8$ ms. The utility function parameters for ED class 1 and 2 are set to $[a_1, b_1] = [1, 10]$ and $[a_2, b_2] = [0.7, 20]$ respectively, unless mentioned otherwise. The penalty parameters $\gamma_1 = \gamma_2 = 1$ unless mentioned otherwise. We initially fix the latency threshold for ED class 1 and 2 to $\delta_1 = b_1 + (4/a_1)$ and $\delta_2 = \infty$ respectively and later study the performance improvement by optimizing δ_1 over $[\min(0, l_{\text{ed},0.99}^1 - t_{\text{av}}^1), l_{\text{ed},0.01}^1]$, where $t_{\text{av}}^1 = 1$ ms in our

simulations. For each simulation scenario, we determine th optimal latency threshold for i^{th} PU class l_t^i by searching ove $[l_d^i - t_{99}^i, l_d^i]$.

A. Impact of delay heterogeneity

Fig. 4 shows plot of system utility for a nearly homogeneou system with $T_{\text{pu}}^1 = T_{\text{pu}}^2 = 0.8$ s, $[l_d^1, l_d^2] = [7.8, 8]$ ms $[a_1, a_2] = [0.65, 7]$ ms^{-1} and $[b_1, b_2] = [19, 20]$ ms. W then increase the heterogeneity in system by using the defaul parameters which results in system utility as shown in Fig. 5 We note that the performance gap between the propose scheduler and the other schemes, ΔV, is much larger for het erogeneous system ($\Delta V = 0.33$) compared to homogeneou system ($\Delta V = 0.19$), especially at higher λ_{ed}^1. This is becaus in the proposed scheduler, the effect of priority order amon classes on the system utility becomes more pronounced wit increasing heterogeneity in classes. Lastly, the performance o proposed scheduler further improves when we drop the failed PU packets due to reduced congestion.

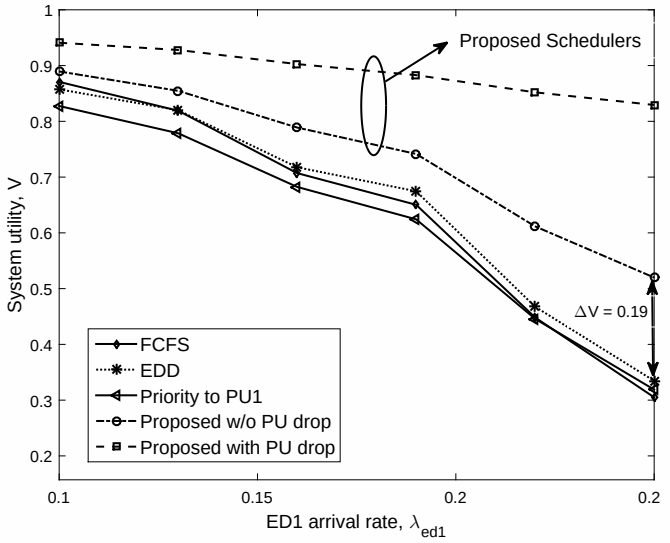

Fig. 4. System utility with delay-homogeneous M2M traffic.

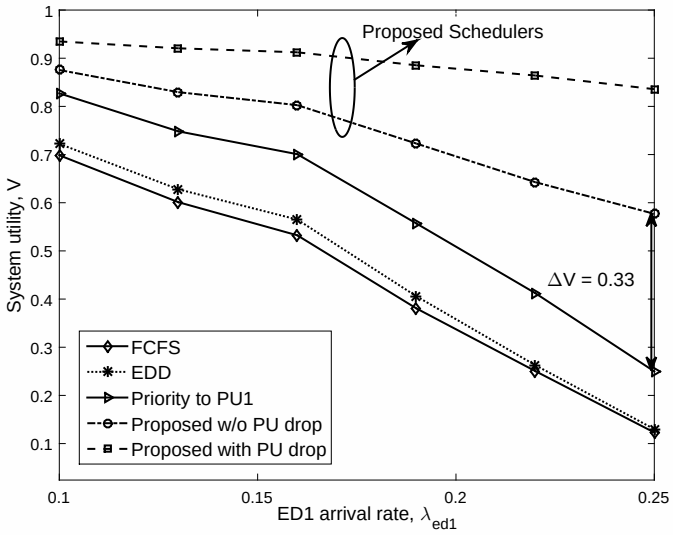

Fig. 5. System utility with delay-heterogeneous M2M traffic.

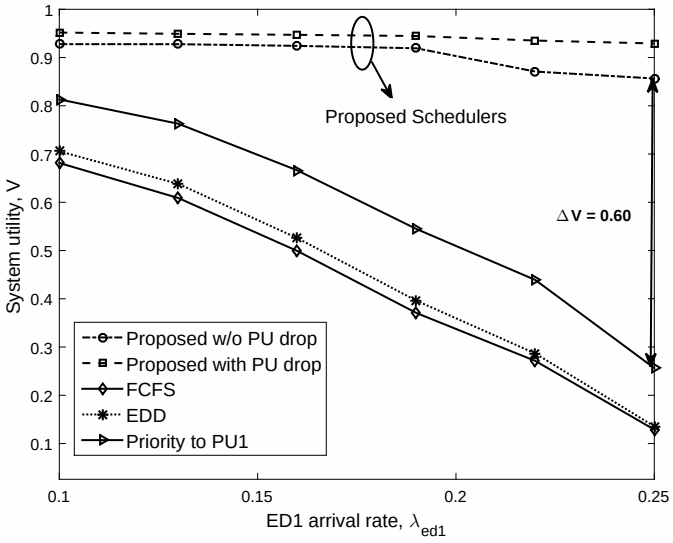

Fig. 6. Impact of optimal latency threshold for ED class 1.

B. Impact of choice of ED threshold

Fig. 6 shows the impact of selecting optimal threshold δ_1 for ED class 1 on the system utility. We note that the performance of proposed scheduler is significantly improved ($\Delta V = 0.60$) compared to Fig. 5 ($\Delta V = 0.33$) wherein δ_1 was arbitrarily set to $b_1 + (4/a_1)$. This is because selecting optimal[2] δ_1 prevents latency of ED class 2 from growing very large when λ_{ed}^1 becomes very large relative to λ_{ed}^2.

C. Impact of penalizing PU failures

We now consider the impact of penalizing PU failures on the performance improvement of different schedulers. Fig. 7 shows the plot of system utility when the penalty factor is small but equal for both PU classes ($\gamma_1 = \gamma_2 = 1.2$). We observe that while the utility for other schedulers take a performance hit and approach 0 at $\lambda_{\text{ed}_1} = 0.25$, the proposed scheduler is fairly robust to the penalization. This is because it possesses multiple degrees of freedom in the form of variable latency thresholds which gets re-calibrated when we begin penalizing PU failures. However, other schedulers lack this flexibility and hence their performance suffer. Further increase in penalty of PU class 1 worsens the performance of other schedulers as shown in Fig. 8. But once again, the proposed scheduler is robust and its performance does not decrease noticeably.

[2]Here we refer to optimality over reduced search space for δ_1.

VI. Conclusions

In this paper, we presented a online delay-efficient multi-class scheduler for delay-heterogeneous M2M uplink traffic. The data from each sensor is classified into PU and ED

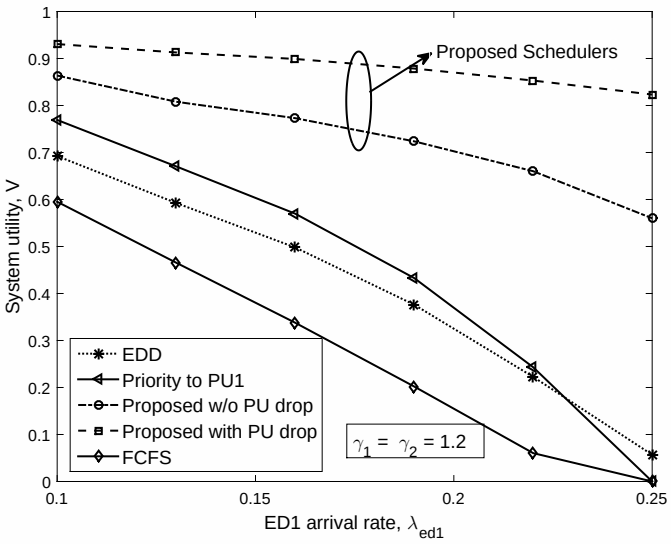

Fig. 7. Utility performance for small but equal PU failure penalty.

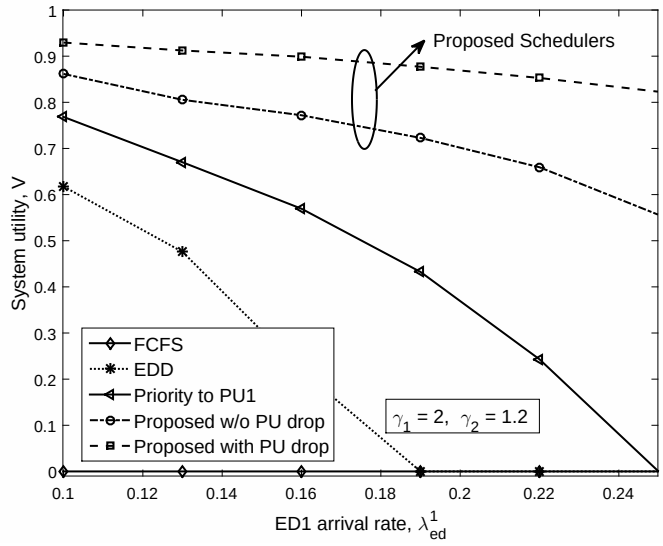

Fig. 8. Utility performance for unequal PU failure penalty.

traffic. Furthermore, due to the delay-heterogeneity across the sensors, the aggregated traffic at M2M application server is classified into multiple PU and multiple ED classes. We use step and sigmoidal functions to represent the different utility of service to PU and ED classes. The average utility of each class is finally aggregated into a proportionally-fair system utility metric. The proposed delay-efficient scheduler uses heuristics that aim to maximize system utility metric. Specifically, it exploits the *firm* utility for PU packets by prioritizing service to ED data as long as we meet the deadline for the PU data. However with increase in network size, an increasingly number of PU packets fail to meet their deadline. We remove the failed PU packets to reduce congestion and improve overall system utility. We also introduce a novel penalty function as part of PU utility to reduce burst PU failures for critical applications. Using extensive simulations, we did a comprehensive delay-performance analysis of the proposed scheduler with respect to other schedulers. We note that the proposed scheduler outperforms other schedulers with the performance gap increasing with heterogeneity in latency requirements and higher penalty for PU failures.

REFERENCES

[1] 3GPP, "Study on RAN Improvements for Machine-type communications," tech. rep., Technical report, TR 37.868, 2012.

[2] A. Gotsis, A. Lioumpas, and A. Alexiou, "Evolution of packet scheduling for machine-type communications over lte: Algorithmic design and performance analysis," in *IEEE Globecom Workshop*, pp. 1620–1625, Dec 2012.

[3] G. Buttazzo, *Hard Real-Time Computing Systems: Predictable Scheduling Algorithms and Applications*. New York: Springer, Third ed., 2011.

[4] A. Kumar, A. Abdelhadi, and T. Clancy, "An *Online* Delay Efficient Packet Scheduler for M2M Traffic in Industrial Automation," in *Submitted to IEEE International Systems Conference (SysCon)*, 2016.

[5] N. Nikaein, M. Laner, K. Zhou, P. Svoboda, D. Drajic, M. Popovic, and S. Krco, "Simple traffic modeling framework for machine type communication," in *International Symposium on Wireless Communication Systems*, pp. 1–5, Aug 2013.

[6] E. Hossain, Z. Han, and H. V. Poor, *Smart Grid Communications and Networking*. Cambridge University Press, 2012.

[7] R. Pratt, "Real-Time Pricing and Demand Response." https://goo.gl/2TwH0c, July 2011.

[8] A. Abdelhadi and T. Clancy, "A Utility Proportional Fairness Approach for Resource Allocation in 4G-LTE," in *IEEE International Conference on Computing, Networking, and Communications (ICNC), CNC Workshop*, 2014.

[9] A. Abdelhadi and T. Clancy, "A Robust Optimal Rate Allocation Algorithm and Pricing Policy for Hybrid Traffic in 4G-LTE," in *IEEE International Symposium on Personal, Indoor, and Mobile Radio Communications (PIMRC)*, 2013.

[10] T. Tia, J. Liu, and M. Shankar, "Algorithms and optimality of scheduling soft aperiodic requests in fixed-priority preemptive systems," *Real-Time Systems*, vol. 10, pp. 23–43, Jan 1996.

[11] L. Schrage, "A Proof of the Optimality of the Shortest Remaining Processing Time Discipline," *Operations Research*, vol. 16, no. 3, pp. 687–690, 1968.